

SciDAC DataGrid Middleware
A High-Performance Data Grid Toolkit:
Enabling Technology for Wide Area Data-Intensive Applications
Quarterly Report October 2004 thru December 2004

Accomplishments this Quarter:

Striping functionality added to GT3.9.2 Development Release

The release of GT3.9.2 marks the first time striping functionality has been in an official release. Early performance testing on the TeraGrid yielded very promising results. Running memory to memory, we were able to achieve 27 Gbs on a 30 Gbs link (90% utilization) with only 32 nodes. This represents nearly perfect linear scaling. Disk to disk results achieved 17.5 Gbs with 64 nodes. We believe the current bottleneck is the disk sub system and intend to contact the sites and the GPFS technical staff to try and improve this performance.

Reliable File Transfer (RFT) Service ported to Web Services Resource Framework (WSRF)

The Reliable File Transfer (RFT) service has been successfully ported from the Open Grid Service Infrastructure (OGSI) to the Web Services Resource Framework (WSRF). We will continue to support the OGSI version at least until the release of 4.2 (late 2005). Initial testing indicates some scalability problems in the WSRF core libraries. We are working with the WSRF core developers to address these issues.

Additional Drivers for the eXtensible Input Output (XIO) System

We completed several new drivers to round out the functionality and usability of XIO. We added a MODE E driver, which enables Open/Close/Read/Write (OCRW) access to the GridFTP data channel (multiple TCP streams). Since MODE E can have out of order arrival of data, we implemented an ordering driver which buffers the blocks and presents them in order to the application. We also implemented a GridFTP driver which allows OCRW access to a file via a GridFTP server. Finally, we implemented a queuing driver to allow multiple outstanding writes. We are also working on a UDT (reliable UDP based protocol from Bob Grossman's group at UIC). Currently it is functional, but we are not getting the performance out of UDT that we should. We are currently working with the UIC crew to resolve these issues.

Replica Location Service Development

During the current quarter, most of our efforts have gone into packaging the latest version of the RLS for inclusion in the next Globus release. This includes improved documentation for the service. We continue to support a growing number of applications that use the RLS in production, including the Laser Interferometer Gravitational Wave Observatory (LIGO), Earth System Grid, the US portions of the CMS and Atlas physics experiments, Nordugrid and others.

Integration of Replica Location Service with POOL

In the current quarter, we succeeded in integrating the Globus RLS with the POOL environment, which provides persistent object management for particle physics applications. The US CMS physics project demonstrated the successful use of the integrated POOL/RLS system for the COBRA application.

Implementation and Evaluation of a Peer-to-Peer Version of the RLS

Continuing our work from the previous quarter, we further developed and evaluated a peer-to-peer version of the Replica Location Service. This system uses a Chord peer-to-peer structured overlay network to distribute RLS mappings among a collection of peer-to-peer Replica Location Index nodes. A paper describing this work was presented at the SC2004 conference in Pittsburgh in November 2004.

Design and Implementation of a WS-RF Data Replicator Service

We have designed and begun the implementation of a data replicator service that generalizes the publication component of the Lightweight Data Replicator system from the University of Wisconsin at Milwaukee. The function of the data replicator service is to ensure that a specified set of files exist on a storage site by comparing the contents of a local replica catalog with a list of desired files, transferring copies of the missing files from other locations and registering them in the local file catalog. The Data Replicator Service (DRS) is a WS-RF service that has interfaces to other Globus services, including the Reliable File Transfer Service and the Replica Location Service. The DRS service will be available as a Technical Preview component in the Globus Toolkit 4.0 beta release.

OGSA DAI Performance Evaluation

We have worked in this quarter on a performance evaluation of the latest release of the OGSA DAI database service. Our goal is to determine whether this version offers sufficient functionality and performance for use as a metadata service for the Earth System Grid project.

Open Science Grid

We are leading the effort on defining data management architecture in the Blueprint document of the Open Science Grid activity. We participated in two Blueprint meetings in the current quarter.

Plans for Next Quarter

A major portion of work on RLS, GridFTP, RFT, and XIO over the next quarter will center on performance testing, hardening, and improvements in ease of use in preparation for the GT4 final release early next year.

We plan to work on moving the Sloan Digital Sky Survey archive from Fermi Lab to Starlight via RFT. This archive contains in excess of 1 million files.

We continue to work on custom data storage interfaces to the High Performance Storage System (HPSS), the Storage Resource Broker (SRB), and the University of Wisconsin Network Storage (NeST) system. This will allow GridFTP clients to access these systems via a GridFTP server.

A Beta release of GT4.0 is planned for December.

We will continue work on the interfaces and implementation for the Lightweight Data Replicator (LDR) scheduled transfer/registration functionality

The current centralized implementation integrates with POOL using a single RLS Local Replica Catalog. In a realistic grid environment, we need to support a distributed RLS deployment. We are currently designing the distributed configuration of the RLS/POOL integration, and we will complete this implementation in the next quarter.